arXiv:1606.00803v1 [cs.NA] 2 Jun 2016

Locality-Aware Laplacian Mesh Smoothing

Guillaume Aupy
Electrical Eng. and Computer Science
Vanderbilt University
Nashville, TN, USA
guillaume.aupy@vanderbilt.edu

JeongHyung Park
Computer Science and Engineering
Penn State University
University Park, PA, USA
jxp975@cse.psu.edu

Padma Raghavan
Electrical Eng. and Computer Science
Vanderbilt University
Nashville, TN, USA
guillaume.aupy@vanderbilt.edu

June 3, 2016

Abstract

In this paper, we propose a novel reordering scheme to improve the performance of a Laplacian Mesh Smoothing (LMS). While the Laplacian smoothing algorithm is well optimized and studied, we show how a simple reordering of the vertices of the mesh can greatly improve the execution time of the smoothing algorithm. The idea of our reordering is based on (i) the postulate that cache misses are a very time consuming part of the execution of LMS, and (ii) the study of the reuse distance patterns of various executions of the LMS algorithm.

Our reordering algorithm is very simple but allows for huge performance improvement. We ran it on a Westmere-EX platform and obtained a speedup of 75 on 32 cores compared to the single core execution without reordering, and a gain in execution of 32% on 32 cores compared to state of the art reordering. Finally, we show that we leave little room for a better ordering by reducing the L2 and L3 cache misses to a bare minimum.

1 Introduction

In HPC systems, when solving partial differential equations using finite element methods for unstructured meshes in parallel on multicore system, high quality meshes are required [8]. This is due to the fact that mesh quality plays a significant role in both the accuracy of the solution produced by a PDE solver, as well as the solver's execution time [16]. However, when attempting to parallelize the mesh smoothing application to obtain solution of PDE efficiently, we do not obtain full performance scalability due to the irregular memory accesses of the application [20].

When doing a computation over a graph, the vertices are stored in a data-structure according to a specific ordering. Strout and Hovland [18] noticed that the data-ordering of irregular HPC applications impacted the performance of these applications. In particular, they showed that a breadth-first search reordering heuristic outperformed all existing ordering heuristics for mesh smoothing [18].

On an identical execution, different orderings have different effects on the reuse distance and on the cache miss rate. The performance is impacted. In Figure 1, we show how different ordering impacts the average reuse distance, the execution time and the cache misses of an execution of LMS. In particular, we compare three different orderings: a random ordering (that unsurprisingly has the worse performance), the original ordering of the mesh (given by the mesh creation algorithm [15]), and the ordering of the `BFS` algorithm introduced by Strout and Hovland [18]. This strongly supports the importance of choosing a good ordering.

In this work, we propose a novel ordering algorithm that outperforms that of Strout and Hovland [18]. We give some insights on the reason of its performance

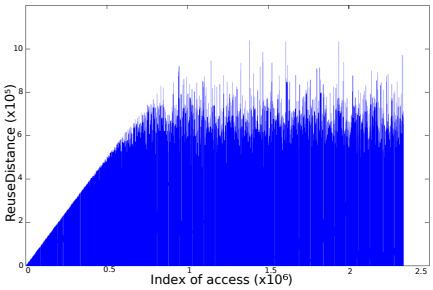

(a) Random ordering, average reuse distance: 90k, L1 cache miss rate: 2.18%, execution time: 10.28s.

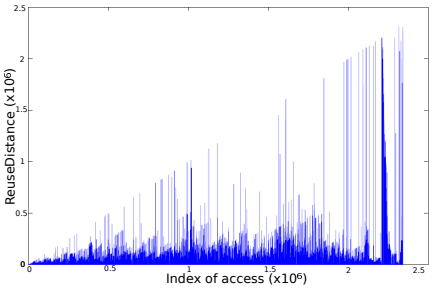

(b) Original ordering, average reuse distance: 4450, L1 cache miss rate: 0.71%, execution time: 7.6s.

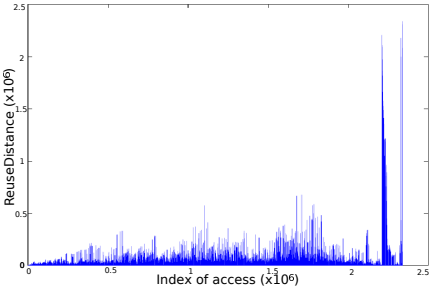

(c) `BFS` ordering, average reuse distance: 2910, L1 cache miss rate: 0.59%, execution time: 6.59s.

Figure 1: Reuse distance plays a significant role in cache performance of HPC applications on multi-core [1]. Simply put, the reuse distance is the number of distinct data accesses between two consecutive accesses of the same data element. We plot here the reuse distance of the first iteration of the LMS algorithm on an *ocean* mesh for different orderings.

by studying thoroughly the reuse distance and cache misses of this algorithm on various meshes. Finally, while our algorithm is cache-oblivious as it focuses in reducing the reuse distance independently of the cache size, we show that it reduces the L2 and L3 cache misses to a bare minimum on the platform we ran it on hence making this algorithm quasi-optimal amongst the possible reordering algorithms.

The rest of this paper is organized as follows. We start by presenting related work on locality-aware iterative algorithms in Section 2. We then introduce the problematic of reuse distance and cache miss (Section 3.1) and the Laplacian Mesh Smoothing algorithm (Section 3.2). Then in Section 4 we describe our novel reordering scheme. Finally, we provide and discuss experimental results in Section 5. Section 6 provides conclusions and direction for future work.

2 Related work

Reordering schemes have proven very efficient to improve performance of many irregular iterative applications such as Feasible Newton mesh optimization [17, 18, 19], mesh warping [14] or Laplacian Mesh Smoothing [11, 17, 18, 19]. In this context, Munson and Hovland [19] developed a Feasible Newton mesh optimization algorithm and benchmark. Their algorithm and benchmark employed both data and iteration reordering in order to improve cache performance. One finding of their research was that reordering of the input data can increase or decrease the number of iterations taken by the inexact Newton method and can affect its success or failure [19]. They compared six different reorderings of the vertices and elements in the mesh. They showed that a breadth-first search of the vertices, followed by a reversing of the order in which the vertices were visited performed better than the original ordering. When data and iteration ordering were performed on the relevant hypergraphs, the reorderings were found to significantly decrease the number of cache misses in all phases of code execution and resulted in significantly faster code [18, 19]. In this work, we compare our results to their results.

Recently, Shontz and Knupp [17] considered mesh vertex reordering techniques to reduce the total time required to improve the mesh quality. Vertex ordering was performed for the first iteration (static) and every iterations (dynamic). Their main finding was that static strategies were superior to dynamic ones because of the overhead of the additional reorderings. Then they compared twenty different orderings from the literature and were not able to find an ordering that stood out as an all-purpose ordering. Park, Knupp and Shontz [11] confirmed this result. Based on their result, this work focuses on an *a priori* ordering. We design a novel ordering that not only stands

out for all meshes considered and outperforms existing heuristics. Unlike this study, we consider cache performance when a reordering technique is applied to mesh smoothing.

3 Problem definitions

In this Section, we introduce the basis for the different notions introduced in this paper.

3.1 Reuse distance and cache misses

Current architectures have different levels of storage systems. They are what are called *hierarchical architectures*: each level has different access costs (as a rule of thumb, the closer the storage system, the faster the access) and different storage sizes (again, often the closest systems have the smallest size).

When a piece of data is accessed, it is stored to the closest level of storage. Then when other data are used, it is pushed back to higher level of storage where the access costs are larger. When a computation tries to access a piece of data and that data is not stored on the closest storage, we call this a *miss*. Then the computation has to go and fetch it from a remote storage which induces additional costs.

In this work we focus on NUMA architectures. In this model, there are four different storage systems: three caches (L1, L2, L3) and a main memory. They are organized with the L1 cache being the smallest and fastest storage system and memory being the largest and slowest storage system.

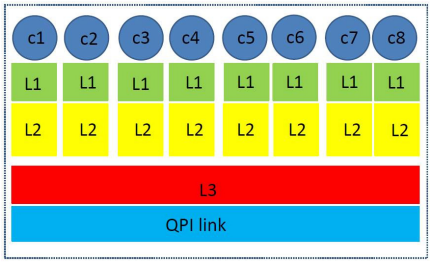

Figure 2: High-level view of a socket of the Intel Westmere-EX processor. The machine has four sockets, which are connected directly via 3.2 GHz QPI links. Each socket has eight cores c_1 to c_8 with the inclusive cache hierarchy; 32K L1 private cache, 256K L2 private cache, and 24MB shared L3 cache.

To choose which unused data is being pushed back to remote storage systems, the caches often follow the Least Recently Used (LRU) algorithm: the piece of data that was the least recently used is the one that will be pushed to remote storage. With this algorithm is introduced the notion of *reuse distance*: between two accesses to the same data, what are the number of distinct piece of data that were accessed.

Intuitively, if the reuse distance is greater than the size of the L1 cache, there will be a L1-cache miss. If it is greater than the size of L2, there will also be an additional L2-cache miss, etc.

In practice the behavior is slightly different [3]. For instance, the fetching is done by cache lines (multiple elements at once) and not by elements which impacts the actual cache misses. However, in this paper we use this simpler theoretical model as a first order approximation to understand and give an intuition of the results observed.

3.2 Laplacian Mesh smoothing

Algorithm 1 *Algorithm for Laplacian Mesh Smoothing*

```
1:  procedure LAPLACIAN SMOOTHING(V, T)
2:      V ← mesh vertex data
3:      T ← mesh triangle data
4:      quality ← 0
5:      for i = 1 to |V| do
6:          compute q_{V[i]}: mesh quality for V[i]
7:          increase quality by q_{V[i]}
8:      end for
9:      Global_quality = (1/|V|) quality
10:     while Global_quality < goal quality do
11:         for i ∈ { interior vertices } do
12:             Smooth V[i] using Equation (1)
13:             Update Global_quality
14:         end for
15:     end while
16: end procedure
```

Mesh smoothing application is performed to improve the quality of the mesh so that an accurate PDE solution can be obtained within a short execution time [4]. In the mesh smoothing procedure, the algorithm first computes the initial mesh quality for a given mesh and do mesh smoothing to improve the mesh quality [12]. After smoothing, we compute the mesh quality again and if the overall mesh quality reached a desirable level, then we stop smoothing. Note that because the desirable quality might never be attained, there is often a maximum number of it-

erations set to the main loop (Algorithm 1, line 10). Note that exact details of the implementation used can be found in the Mesquite software package [2].

We use edge-length ratio [7], i.e., the ratio of minimum and maximum length edges, as a mesh quality metric for computing the mesh quality in this study. The mesh quality for each vertex can be represented as an average quality metric value of triangles that are attached on the vertex. The mesh quality for the entire region of the mesh can be computed by averaging all mesh quality values obtained from each vertex. The range of edge-length ratio mesh quality values is $0 \sim 1$. If the quality value for a triangle is close to 1, we can say the triangle has a good shape, i.e., is close to an equilateral triangle. The goal of the mesh smoothing algorithm is to maximize the average quality values for each vertex.

To improve the quality of the mesh, we perform Laplacian smoothing to replace a vertex position using neighbor vertices coordinates. Suppose there is a vertex v we want to move and N neighbored vertices surrounding the vertex. If we represent the position for i^{th} neighbored vertex as p_i, the new position for $\overline{p}_v$ will be

$$\overline{p}_v = \frac{1}{N} \sum_{i=1}^{N} p_i. \quad (1)$$

Figure 3 shows the initial mesh and the output mesh of Laplacian smoothing.

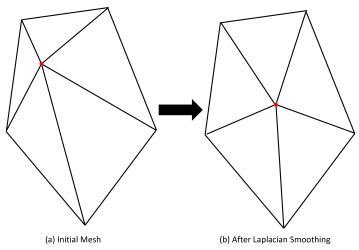

Figure 3: Laplacian smoothing performed on initial mesh. The vertex position inside the mesh was changed.

We would like to improve the performance of the Laplacian mesh smoothing by reordering the initial mesh. In order to do so, we reorder each vertex based on the initial mesh quality for each vertex. Details are described in Section 4.

4 Mesh Reordering to Improve Mesh Smoothing Performance

4.1 Factors Affecting Temporal and Spatial Locality

There are two factors that affect the performance of Laplacian mesh smoothing [18].

Spatial locality is defined as the reuse within a cache line. The nodes of a mesh are stored in the memory. When a node is selected, the node is streamed to the cache along with its neighboring nodes (as many as can fit in a cache line). *Temporal locality* is defined as the reuse of a node that is already stored in the cache before it's cache line is discarded.

In both cases, the absence of sufficient data locality causes a cache miss, thus adversely impacting the execution time.

Since the Laplacian smoothing method processes a node and some of its neighbors successively, prefetching the neighboring nodes into the caches improves the application's cache performance. In particular this is designated to improve the temporal locality: since processing a node requires information from its neighbors, it seems natural to process its neighbors then while the information is still cached. A good reordering strategy then aims at improving the spatial locality.

Intuitively, when one executes a node, one needs to fetch its neighbors, hence having them close by together should improve spatial locality. In this context, while naive, `BFS` seems to be a good start to improve spatial locality.

Figure 4 shows two partial traces of node visiting observed from Laplacian mesh smoothing when two different reuse distances are applied. When the reuse distances for Laplacian mesh smoothing are reduced by applying `BFS` ordering, temporal and spatial localities are significantly improved. This example dictates the interplay between temporal and spatial localities on one hand and reducing reuse distances on the other.

4.2 Toward a Reuse Distance Reducing Ordering

We now would like to consider how to reduce the reuse distances for Laplacian mesh smoothing.

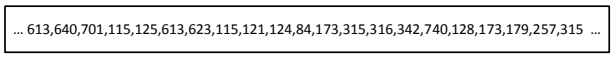

(a) Traces for Laplacian mesh smoothing with DFS ordering.

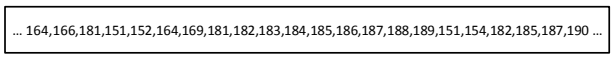

(b) Traces for Laplacian mesh smoothing with BFS ordering

Figure 4: Partial traces observed from node visiting for Laplacian mesh smoothing. The number represent the location of the data accessed in the data array. The closer the numbers, the shorter the distance between the accesses.

The main idea of doing a reordering of the vertices is to improve locality when they are accessed by the algorithm. By the time the Laplacian mesh smoothing terminates, there is a certain order in which nodes have been considered.

We give an example of the usefulness of a good ordering in Figure 5. Consider the synthetic mesh shown in Figure 5. Let us consider two orderings: the Depth First Search ordering given in Figure 5a and the Breadth First Search ordering given in Figure 5b. Assume that node j (10 in DFS, 3 in BFS) has worse quality. Then during the execution of the algorithm (Read Data array), it will be accessed first, followed by its neighbors: k, m, i, a, b (to compute its Laplacian value). In the DFS ordering, the range of accesses in the node array varies between positions 1 and 10, while in the BFS ordering, it varies between positions 1 and 7. Minimizing the span of accesses allows for a better spatial locality.

The nodes in the mesh reordered by BFS ordering store their nodes spatially. This causes the access patterns for Laplacian mesh smoothing to become similar to the memory access of nodes streamed in the cache. Though the accessed node list becomes similar to the streamed node list, there are still spaces for improvement for obtaining optimal reuse distance of Laplacian mesh smoothing.

We now consider improving temporal locality for Laplacian mesh smoothing.

The Laplacian mesh smoothing is a greedy algorithm. When smoothing the mesh, the LMS algorithm starts by visiting the node that has the worse quality. Once the smoothing process for the node is over, it selects another node that has the worst quality among nodes nearby the node, i.e., neighboring nodes.

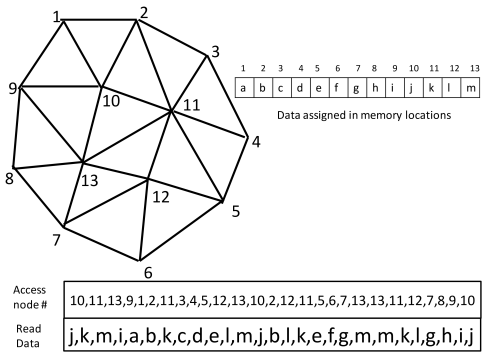
1 2 3 4 5 6 7 8 9 10 11 12 13
a b c d e f g h i j k l m
Data assigned in memory locations
Access node #
10,11,13,9,1,2,11,3,4,5,12,13,10,2,12,11,5,6,7,13,13,11,12,7,8,9,10
Read Data
j,k,m,i,a,b,k,c,d,e,l,m,j,b,l,k,e,f,g,m,m,k,l,g,h,i,j

(a) DFS ordered Mesh.

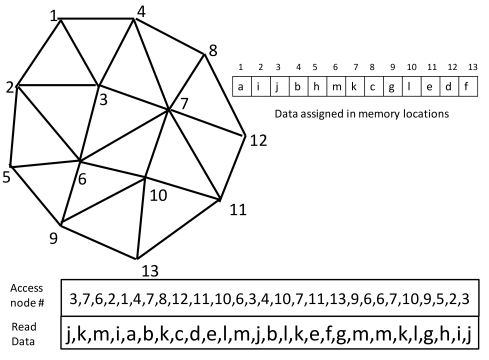
1 2 3 4 5 6 7 8 9 10 11 12 13
a i j b h m k c g l e d f
Data assigned in memory locations
Access node #
3,7,6,2,1,4,7,8,12,11,10,6,3,4,10,7,11,13,9,6,6,7,10,9,5,2,3
Read Data
j,k,m,i,a,b,k,c,d,e,l,m,j,b,l,k,e,f,g,m,m,k,l,g,h,i,j

(b) BFS ordered Mesh.

Figure 5: For both mesh, the data values (a to m) of each nodes are stored in a 13 consecutive memory locations. While the ordering changes the location of the different data values in the data array, the smoothing algorithm is identical (see the sequence of read data).

We studied reuse distance patterns of the different iterations of the mesh smoothing algorithm. As can be seen on Figure 6 (note that we had very similar results for other meshes), the reuse distance has similar patterns over the different iterations (there are eight iterations in the execution plotted on Figure 6). We conjecture that the access patterns for Laplacian mesh smoothing can be controlled by the initial qualities of each node in the mesh.

Hence we conjecture that if we sort the nodes and their neighboring nodes based on the qualities they have, the temporal locality will be improved. We base our reordering heuristic on this conjecture.

Under this conjecture, we propose a mesh reordering scheme called `RDR` in order to reduce the reuse distance of Laplacian mesh smoothing. Our reordering scheme (Algorithm 2) follows a similar pattern to

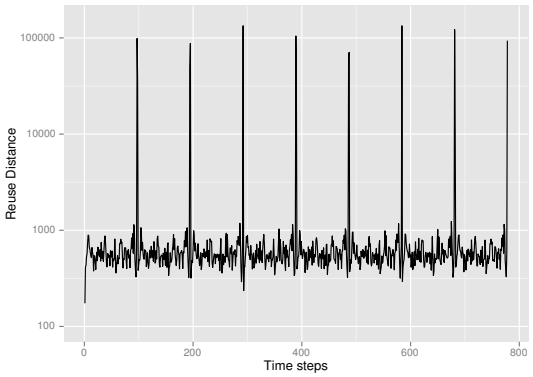

Figure 6: Observed reuse distance profiles for a *carabiner* mesh given the initial ordering. In order to make it more readable, we divided each iteration into 100 *Time Steps* where each time step is the average of approximatively 20,000 consecutive data accesses.

the mesh smoothing algorithm iterations. Starting from the node with the worse quality,

1. From a given node already ordered, we sort all its neighbors that have not been ordered yet by increasing quality.
2. We append them as such to the list of already ordered neighbors.
3. Then we perform the same algorithm for its neighbor of worse quality (that has not been processed yet).

Theorem 1. *Given a mesh (V, T), then Algorithm 2 orders all elements of the mesh exactly once.*

Proof. The fact that each element of the mesh is ordered at most once is guaranteed by the array `sorted`: once element $V[i]$ is added to $V_{\texttt{new}}$, then the value `sorted`$[i]$ is set to True and never set back to False. Furthermore, it can only be added to $V_{\texttt{new}}$ if `sorted`$[i]$ = False.

The fact that each element is ordered at least once is guaranteed by the property that (i) at the end of the algorithm, $\forall i \in$ {interior vertices}, `processed`$[i]$ = True (line 7 guarantees it), and (ii) `processed`$[i] \implies$ `sorted`$[i]$ (every time `processed`$[i]$ is set to True, `sorted`$[i]$ was also set to True earlier (line 12→line 10, line 22→line 19)). □

When mesh smoothing application is executed, the application calculates the initial qualities of each vertex in a mesh, smooths the mesh vertices, and then

Algorithm 2 *Algorithm for Reuse Distance Reducing Ordering*

```
1: procedure RDR(V, T)
2:     V_new ← empty array of size |V|
3:     processed ← bool array of size |V| initialized with false
4:     sorted ← bool array of size |V| initialized with false
5:     next_num ← 1
6:     for i ∈ {interior vertices sorted by increasing quality} do
7:         if not processed[i] then
8:             if not sorted[i] then
9:                 V_new[next_num] ← V[i]; incr next_num
10:                sorted[i] ← True
11:            end if
12:            processed[i] ← True
13:            l ← {unprocessed neighbors of i sorted by increasing quality}
14:            while l ≠ ∅ do
15:                for j = 1 to |l| do
16:                    if not sorted[l[j]] then
17:                        V_new[next_num] ← V[l[j]]
18:                        incr next_num
19:                        sorted[l[j]] ← True
20:                    end if
21:                end for
22:                processed[l[0]] ← True
23:                l ← {unprocessed neighbors of l[0] sorted by increasing quality}
24:            end while
25:        end if
26:    end for
27:    return V_new
28: end procedure
```

computes the quality of the mesh to see the difference between initial and final quality of the mesh. We find that a vertex which has a bad quality in the beginning requires more time to be smoothed compared to other vertices. If such vertices are processed earlier than other vertices, the neighbors of the vertex which has the worst quality are already in the cache and thus, we can improve the spatial locality. The neighboring vertices are also reordered in increasing order.

By reordering all the vertices based on the methodology we described above, we are able to obtain a node trace which is very similar to the initial streamed node list for Laplacian mesh smoothing. Hence, we make sure that both temporal and spatial localities are improved.

5 Experimental Results

In this section, we evaluate our Reuse Distance Reducing ordering (`RDR`) for improving the performance of Laplacian mesh smoothing. We first describe the experimental setup used in this study. We then show the execution times for Laplacian mesh smoothing reduced by our reuse distance reducing ordering. We also test the scalability of the Laplacian mesh smoothing with our reuse distance reducing ordering.

5.1 Experimental Setup

System Setup. We used an Intel Westmere-EX architecture system to evaluate our reuse distance reducing ordering. The Intel Westmere-EX architecture is equipped with 4 eight-core Intel Xeon E7-8837 processors. It supports 32 concurrent threads. Each core has 32K L1 private cache and 256K L2 private cache with reported access latencies of 4 and 10 cycles respectively [9], and they share 24MB L3 cache. The L3 cache serves as the central unit for on-chip inter-core accesses and accesses to off-chip processors include latencies of data transfer on the QPI link. Consequently, L3 data access latencies can vary from 38 to 170 cycles depending on core location and cache-line state [9]. The machine is an inclusive cache hierarchy. Each processor is directly connected to other three processors via 3.2 GHz QPI links. Additionally, access to the memory can range from 175 to 290 cycles [9]. Figure 2 shows the high-level view of the Intel Westmere-EX processor.

For multi-thread running, OpenMP library is used for both systems. Thread affinity is set via KMP_AFFINITY=compact, granularity=fine for pinning each thread to each core. In all cases, thread scheduling is set to be static for simply collecting the application trace of each thread by evenly dividing the vertices. We implemented parallel Laplacian mesh smoothing based on the module in Mesquite [2]. For the purpose of this evaluation we put a quality convergence criterion to 0.000005 (meaning if the quality has improved by less than this criterion, the execution stops). **Note that the orderings did not change the number of iterations needed to reach this criterion**

Label	Mesh	vertex	triangle
M1	carabiner	328082	652920
M2	crake	298898	595638
M3	dialog	306824	611620
M4	lake	375288	747676
M5	riverflow	332699	661615
M6	ocean	392674	783040
M7	stress	312763	622868
M8	valve	300985	599368
M9	wrench	386757	771097

Table 1: Input Mesh Configuration

Test Suite. To determine the impact of reuse distance reducing ordering on the mesh smoothing process, we tested nine meshes shown in Figure 7 (Coarse approximations are shown). The meshes were generated by Triangle [15] and Table 1 gives their configurations.

5.2 Performance: Execution Time and Reduced Reuse Distance

5.2.1 Serial execution

We first test our reuse distance reducing ordering on a serial run of Laplacian mesh smoothing. Figure 8 shows the execution time results of Laplacian mesh smoothing when `RDR` was applied. For evaluation purpose, we provide the original (`ORI`) reordering (as computed by Triangle) and Breath First Search (`BFS`) reordering (developed by Strout et al. [18]) together. On the nine meshes, our algorithm got a 1.39 average speedup compared to `ORI` and 1.19 speedup compared to `BFS`.

5.2.2 Cache Performance

We collected the cache performance counter using PAPI 5.1.0.2 [10] to compute cache performance. Figures 9a, 9b and 9c show the L1, L2, and L3 cache performance for Laplacian mesh smoothing running on a

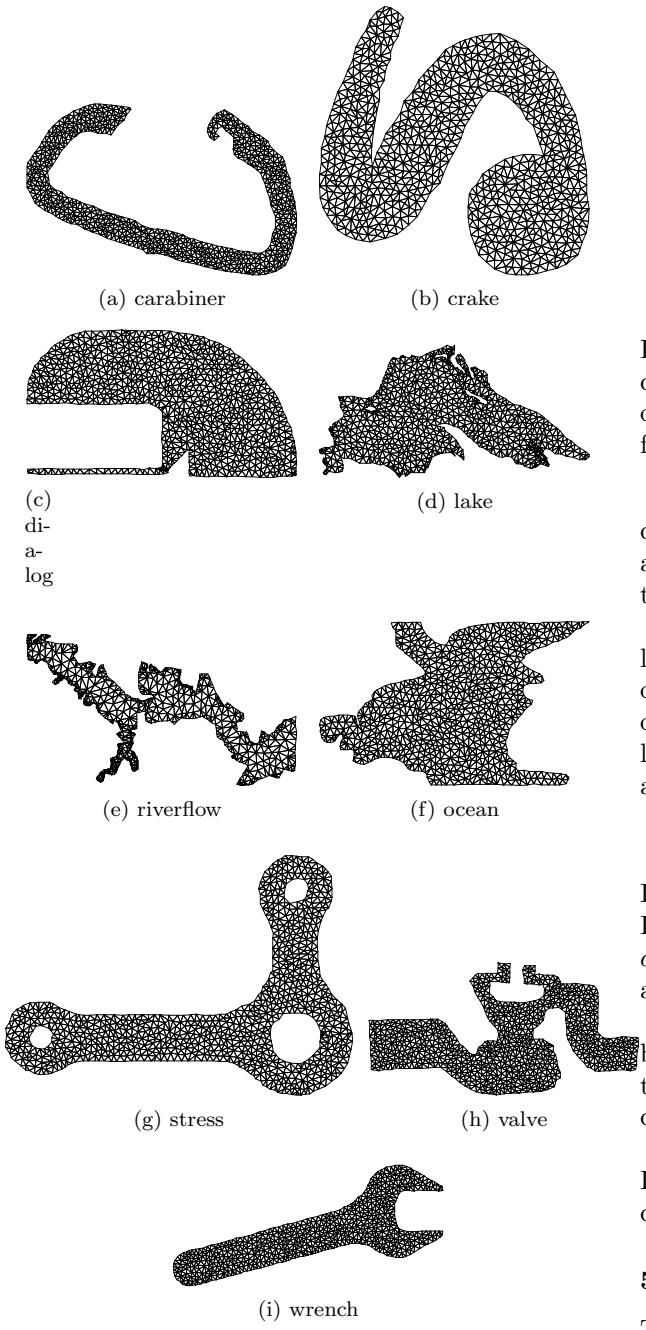

(a) carabiner (b) crake

(c) dialog (d) lake

(e) riverflow (f) ocean

(g) stress (h) valve

(i) wrench

Figure 7: 2D meshes generated by Triangle [15] and used in the experiments (coarser but representative versions).

single core for the different orderings. After the `RDR` reordering was applied to Laplacian mesh smoothing, cache miss rates significantly decreased. On average on one core, there were 25% (resp. 6.3%) fewer L1 cache misses, 71% (resp. 51%) fewer L2 cache misses and 84% (resp. 65%) fewer L3 cache misses compared to `ORI` (resp. `BFS`).

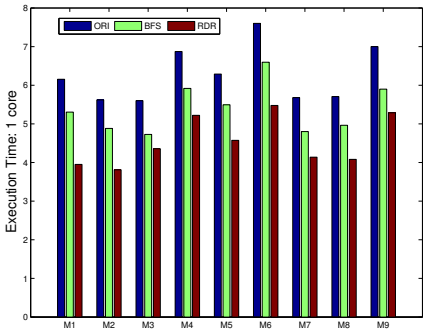

Figure 8: Execution time (seconds) results for Laplacian mesh smoothing when reuse distance reducing ordering was applied. `RDR` ordering is 1.39 times faster than `ORI` ordering.

To understand how this affects the execution time, let us call m_1 (resp. m_2, m_3) the cache miss rates of L1 (resp. L2, L3), and c_2 (resp. c_3, c_m) the cost of a cache access to L2 (resp. L3, Memory). Then let #accesses be the number of data accesses, the additional cost to the execution time is:

$$(m_1 \cdot c_2 + m_1 m_2 \cdot c_3 + m_1 m_2 m_3 \cdot c_m) \cdot \#\text{accesses}. \quad (2)$$

Indeed, $m_1 \cdot \#$accesses operations are not found in L1 and are fetched in L2 (hence an additional cost of c_2), out of those, $m_2 \cdot \#$accesses are not found in L2 and have to be fetched in L3 etc.

To give an example of what it represents, for Carabiner, with the access costs of the machine (see Section 5.1), `ori` has 927k additional clock cycles due to cache misses, `bfs` has 528k, and `rdr` has 210k.

Similar cache performance is obtained when the Laplacian mesh smoothing runs on multicores (32 cores).

5.2.3 Reuse Distance study

To understand the speedups observed during the serial execution, we study the reuse distance of the different orderings. To measure the reuse distance, because this functionality is not offered by PAPI, we did a verbose run noting the data locations being addressed and analyzed them.

We present some of the quantiles observed in Table 2. The X quantile is defined as the smallest value such that there at least a proportion X of the population below it.

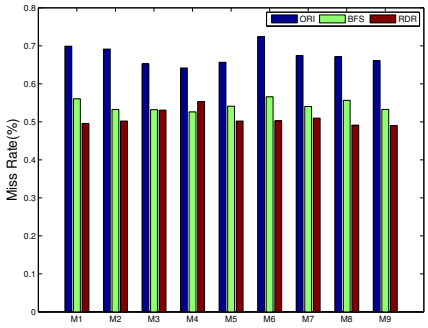

(a) L1

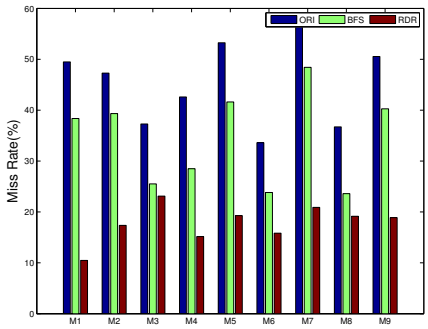

(b) L2

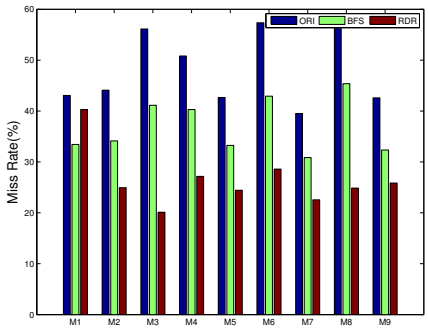

(c) L3

Figure 9: Cache performance results on one core when reordering were applied. The LX miss rate depends on the number of LX accesses: a higher L2 (resp. L3) miss rate does not necessarily means higher number of misses, they need to be put in relation with the number of L1 (resp. L1 and L2) misses.

What is notable is that for most of the data-accesses, the reuse distance is well below that of the L1 cache: as an order of magnitude, assuming that each node is 66 bytes[1], and that we use the theoretical cache miss model (Section 3.1); then below a reuse distance of 496 (resp. 3970; 372,000) there should not be any L1 (resp. L2; L3) cache miss. This is consistent with the findings that the L1-cache miss rate is very low for all orderings (see Figure 9).

It is however important to notice that the maximum reuse distance of `RDR` is well below the theoretical reuse distance for a L3 cache miss (at least 42 times smaller). Intuitively, even if our approximation is only a first order approximation, this means that none of the L3 cache misses observed by PAPI in `RDR` are due to elements being reused. Most certainly they are either compulsory cache misses or due to the first fetching of a given element.

Based on the measured cache miss rates, we have been able to estimate the actual number of misses which we represent in Table 3 (Estim. number of misses). We have seen in Table 2 that in practice the L3 cache misses observed in `RDR` were not results of the reuse distance but of other factors, hence we could compute them for all graphs[2] and subtract them from the total number of L1, L2 and L3 misses for all orderings.[3] The results obtained allow us to go further in the understanding of the behavior of the different algorithms.

In order to do this, we again used the theoretical model introduced in Section 3.1 to measure the maximum number of elements that fit in the cache. We formulate this model:
Assuming that there are n_X LX misses, then the n_X accesses with the largest reuse distance are the one that missed.

With this approximation we measured the number of elements that could fit in the cache for each graph and each ordering. We report those values in Table 3 (section Estim. max number of elements). The first observation is that in general, given a graph, for both orderings `ORI` and `BFS`, the estimated maximum number of elements that fit the different cache are similar (same order of magnitude). This makes sense as for

[1]A node is characterized by (i) its coordinates: two double precision floats, (ii) its connectivity: an array of long integers (typically here five or six neighbors) and (iii) a fixed/boundary state: an integer that says whether or not the vertex is an interior one [2]. Note that this is an approximation as other factors have to be accounted for (for instance the data-structure used to store the node information), in practice the size can be many more times this.

[2]In practice, when computing this number, we were able to verify that it was almost constant for all graphs (between 400 and 500), hence confirming the intuition that it was due to external factors.

[3]Results were similar when we did not perform this operation.

		Estim. number of misses (x10^3)			Estim. max number of elements (x10^3)		
mesh	Ordering	L1	L2	L3	L1	L2	L3
carabiner	ORI	106	50.6	19.9	13.2	21.3	330
	BFS	84	30.2	7.94	10.2	21.2	1060
	RDR	73.9	4.82	0	1.6	1.88	1.94
crake	ORI	95.3	43.4	17.4	24.6	40.9	198
	BFS	72.6	26.7	7.07	18.3	39.2	986
	RDR	68.3	9.3	0	3.4	3.77	3.9
dialog	ORI	91.7	31.9	16.3	59	87.7	108
	BFS	74.1	16.2	4.55	53.2	89.3	157
	RDR	74	14.3	0	5.84	6.05	6.2
lake	ORI	111	44.7	20.7	62.2	95.9	143
	BFS	89.8	22.7	6.72	56.1	96.1	219
	RDR	94.7	10.9	0	8.42	8.65	8.77
riverflow	ORI	99.8	51.4	19.8	30.4	66.8	863
	BFS	81.6	31.8	8.08	20.3	67.6	1290
	RDR	75.4	11.5	0	4.73	5.27	5.43
ocean	ORI	131	41.3	21.9	63.3	300	1050
	BFS	102	21	6.57	41.6	770	1830
	RDR	89.8	10.6	0	5.83	5.99	6.12
stress	ORI	96.6	54.6	19.4	71.5	91.7	130
	BFS	76.7	35.3	8.42	59.1	87.6	137
	RDR	72.2	12.3	0	4.36	4.66	4.81
valve	ORI	92.6	31.9	16.9	37.3	77.6	186
	BFS	76.3	15.4	5.17	28.6	110	1040
	RDR	67	10.1	0	6.68	7.17	7.26
wrench	ORI	117	57.1	21.8	75.3	105	148
	BFS	93.8	35.1	8.38	62.9	103	170
	RDR	86	12.7	0	5.69	6.31	6.49

Table 3: Estimated number of misses for the different levels of cache. Based on those, we estimated the maximum number of elements (using reuse distance) that could fit in the caches.

a given graph the elements have the same size[4] –not necessarily 66 bytes which was an approximation–.

This is particularly visible for the L2 cache. L1 is smaller, hence it is more sensible to the cache optimizations and hence to the first order approximation. On the contrary, the number of accesses in L3 are smaller, hence they are more sensible to the noise incurred by other elements of the computation that are not reported by PAPI. The number of accesses of L2 and its size are the right balance to verify our approximation.

The second observation, is that the estimated number of elements that fit the L2 cache for RDR is many orders of magnitude below that of the other orderings. On the contrary it is very close to that of the L1 cache miss of RDR. Intuitively, this tends to give the idea that there are actually no L2 cache miss with the ordering, and that the L2 cache misses are, similarly to the L3 cache misses linked to other factors (for this study we only take into account the reuse distance).

This would show that the RDR ordering not only does not have any L3 cache misses, but also have very few L2 cache misses, making it quasi-optimal.

5.3 Performance Scalability

Finally, we have studied the scalability of the reordering. In order to do so we have run the Laplacian mesh smoothing algorithm on 1, 2, 4, 8, 16, 24 and 32 cores.

In Figure 10, we compare the different speedups per mesh. Each speedup is relative to a serial baseline (the execution time of ORI on one core), hence given by the formula:

$$\text{Speedup}(\texttt{ordering}, p) = \frac{T_{\texttt{ORI}}(1)}{T_{\texttt{ordering}}(p)}$$

The first notable result is that even with only the original ordering (ORI) the speedup is supra-linear. One way to explain this result would be a better data management when using multiple cores. We plot the number of additional accesses for the three first meshes Carabiner, Crake and Dialog in Figure 11. Results are similar for the remaining meshes. We can observe that with the scalability, the distance with respect of the data accessed decreases. This could explain part of the superlinear speedup observed, in particular, we can see that the fluctuations of the slopes of the speedup of both Crake and Dialog seem to follow the fluctuations of the number of memory accesses slope. When we compare the speedup per ordering (that is: $T_{ord}(1)/T_{ord}(p)$), both BFS and RDR also have the same superlinear speedup than ORI. Hence we suspect that this superlinear speedup has more to do with the architecture (for instance with less than four threads, they might be distributed in a "scattered" way, leading to four times the L3 caches from one to four cores) or the LMS than with the reorderings (or lack of). The fact that (i) this observation is independent of the ordering, and (ii) the speedup becomes more "normal" from 4 to 32 cores (about 7) increases the likeliness of this hypothesis.

The second notable result is that the speed-up gained by our algorithm is greater than the BFS ordering on almost all data sets. Furthermore, the average speed-up is clearly dominant as can be seen in Figure 12, with a speed-up greater than 75 when RDR is applied on 32 cores!

Finally, we plot on Figure 13 the gain in execution time of RDR compared to the execution time of

[4]Note that this is not as true for ocean and valve. We have no explanation for these graphs.

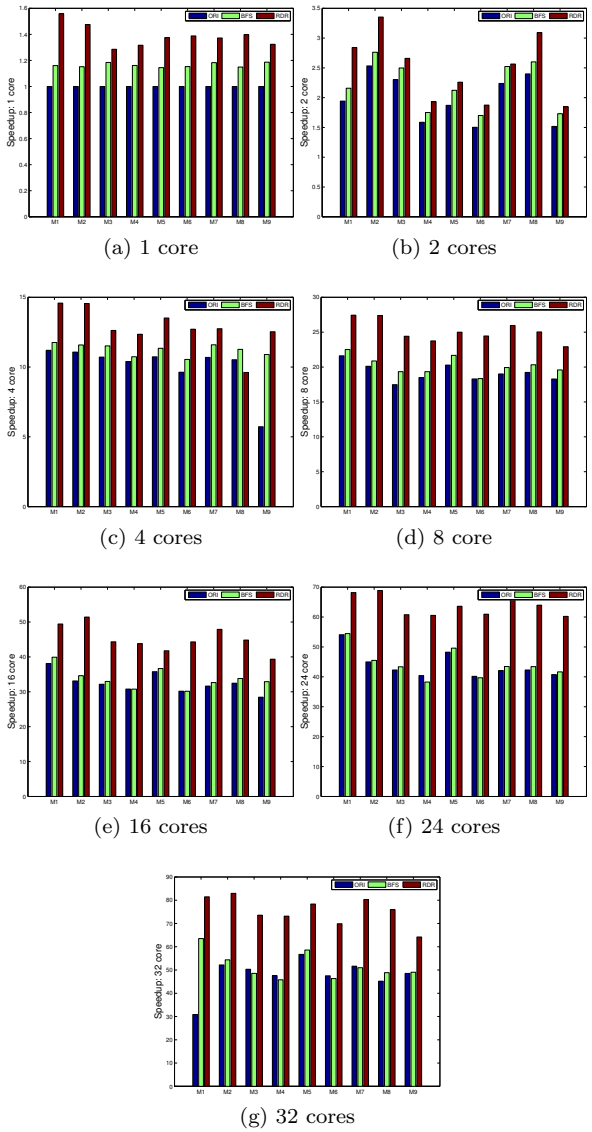

(a) 1 core (b) 2 cores

(c) 4 cores (d) 8 core

(e) 16 cores (f) 24 cores

(g) 32 cores

Figure 10: Observed speedup relative to the serial `ORI` baseline. Reuse distance reducing ordering (`RDR`) provides significant performance improvement.

`ORI` and `BFS`. Clearly `RDR` is dominant on most of the meshes and not only on average. The only exception is the execution of Valve on 4 cores. Overall, the expected gain over the `ORI` ordering varies between 20% and 30% depending on the number of cores, and between 10% and 30% compared to the `BFS` reordering.

5.4 Discussion on reordering cost

To conclude this experimental evaluation, we discuss the additional cost incurred by the pre-computation

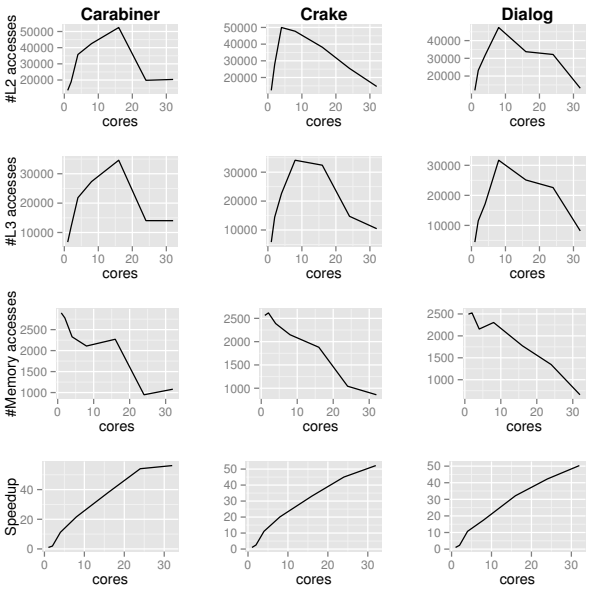

Figure 11: Number of accesses as a function of the number of cores used for `ORI` for different meshes. Overall the distance where the data is fetched decreases with the number of cores.

(reordering). Because it follows so closely the actual mesh smoothing algorithm, our reordering has a cost of approximatively one iteration with the `ORI` ordering.

Hence, with an average gain between 20 and 30% depending on the number of cores (Figure 13) over the computation with the `ORI` ordering, it follows that the gain will be notable as soon as there are more than four iterations. To conclude, we would not recommend doing any reordering when the initial mesh quality is close to the desired mesh quality as one should not expect many iterations. However if it is

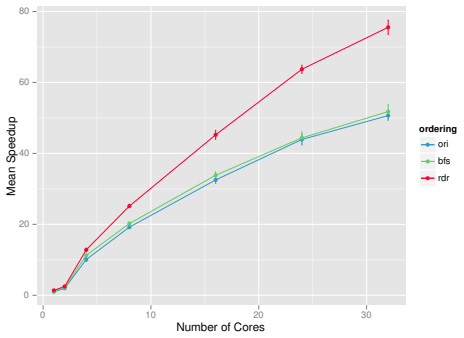

Figure 12: Mean speedup versus $T_{\texttt{ORI}}(1)$.

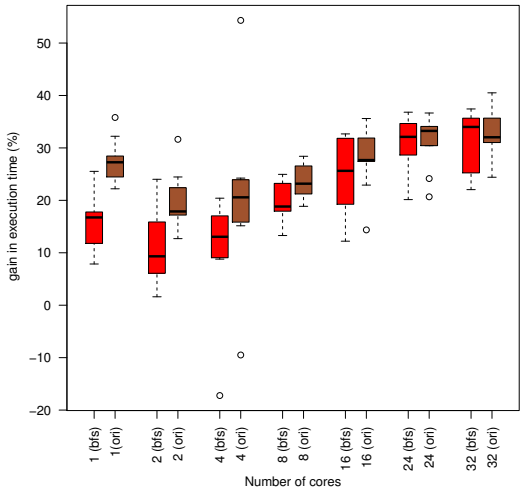

Figure 13: Gain with scalability the performance gain is $\frac{T_{\texttt{algo}}(x)-T_{\texttt{RDR}}(x)}{T_{\texttt{algo}}(x)}$, for `algo` being either `ORI` of `BFS` and x being the number of cores.

not the case, one should clearly use the reordering we designed.

6 Conclusion

In this work, we have presented a pre-computation heuristic for Laplacian mesh smoothing. Our pre-computation takes the form of a reordering of the initial data and performs remarkably well compared to a state of the art reordering heuristic.

Our reordering is based on an observation we made, the reuse distance of the mesh smoothing algorithm do not vary much over iterations. We hypothesized that a good order for the first iteration would be efficient for subsequent iterations. We have developed and evaluate `RDR`, a scheme that takes into account the initial qualities of each vertex and reorders the mesh based on the quality in ascending order to improve both temporal and spatial localities for memory accesses of Laplacian mesh smoothing.

Thanks to our reordering, we decrease the number of cache misses of the Laplacian Mesh Smoothing by 25% (L1), 71% (L2) and 84% (L3) on average on a single core. Similar cache performance is obtained when we run the application on multicores (up to 32 cores). In turn, this decrease in cache miss allowed us to reach a mean speed-up of 75 when running on 32 cores compared to the single core performance with no reordering, and a gain in execution time of 30% compared to the 32-cores `ORI` execution with as little as eight iterations of the mesh smoothing algorithm. We were able to justify those gains by showing that our algorithm is quasi-optimal with regard to the number of L2 and L3 misses.

By modifying almost nothing in the original algorithm, we have shown how much cache-misses impact the execution time of the Laplacian mesh smoothing algorithm. We expect our new reuse-distance-aware algorithm to outperform extensions of Laplacian mesh smoothing as well. We conjecture that either this ordering or an ordering based on the idea that it needs to be efficient for the first iteration, could improve other mesh application performances such as mesh untangling [6], constraint mesh smoothing [13], and mesh swapping [5].

Finally, this result shows that data-locality is critical in the execution of applications, and that a short pre-computation may improve drastically the execution of an application.

Acknowledgements

Part of this work was done while the authors were at the Pennsylvania State University. This material is based upon work supported by Vanderbilt University and the National Science Foundation, CCF #1319448.

References

[1] K. Beyls and E. D'Hollander. Reuse distance as a metric for cache behavior. In *Proceedings of the Parallel and Distributed Computing and System*, pages 617–662, 2001.

[2] M. Brewer, L. Freitag, P. Knupp, T. Leurent, and D. Melander. The Mesquite Mesh Quality Improvement Toolkit. In *Proc. of the* 12^{th} *International Meshing Roundtable*, pages 239–250. Sandia National Laboratories, 2003.

[3] U. Drepper. What every programmer should know about memory. *Red Hat, Inc*, 11, 2007.

[4] L. Freitag, P. Knupp, T. Munson, and S. Shontz. A comparison of optimization software for mesh shape-quality improvement problems. In *Proc. of the* 11^{th} *International Meshing Roundtable*, pages 29–40. Sandia National Laboratories, 2002.

[5] L. Freitag and C. Ollivier. Tetrahedral mesh improvement using swapping and smoothing. *Int. J. Num. Meth. Eng.*, 40:3979–4002, 1997.

[6] L. Freitag and P. Plassmann. Local optimization-based simplicial mesh untangling and improvement. *Int. J. Numer. Math. Eng.*, 49:109–125, 2000.

[7] P. M. Knupp. Algebraic mesh quality metrics. *SIAM J. Sci. Comput.*, 23(1):193–218, Jan. 2001.

[8] J. Mellor-Crummey, D. Whalley, and K. Kennedy. Improving memory hierarchy performance for irregular applications using data and computation reorderings. *Int. J. Parallel Program.*, 29(3):217–247, June 2001.

[9] D. Molka, D. Hackenberg, R. Schöne, and M. S. Müller. Memory performance and cache coherency effects on an intel nehalem multiprocessor system. In *18th International Conference on Parallel Architectures and Compilation Techniques. PACT'09.*, pages 261–270. IEEE, 2009.

[10] P. J. Mucci, S. Browne, C. Deane, and G. Ho. Papi: A portable interface to hardware performance counters. In *Proceedings of the Department of Defense HPCMP Users Group Conference*, pages 7–10, 1999.

[11] J. Park, P. Knupp, and S. Shontz. Static vertex reordering schemes for local mesh quality improvement. In *Technical report, Sandia National Laboratories*, 2010.

[12] J. Park and S. M. Shontz. An alternating mesh quality metric scheme for efficient mesh quality improvement. *Procedia Computer Science*, 4(0):292 – 301, 2011.

[13] V. Parthasarathy and S. Kodiyalam. A constrained optimization approach to finite element mesh smoothing. *Finite Elements in Analysis and Design*, 9:309–320, 1991.

[14] S. Sastry, E. Kultursay, S. Shontz, and M. Kandemir. Improved cache utilization and preconditioner efficiency through use of a space-filling curve mesh element- and vertex-reordering technique. *Engineering with Computers*, 30(4):535–547, 2014.

[15] J. Shewchuk. Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator. In *Applied Computational Geometry: Towards Geometric Engineering*, volume 1148, pages 203–222. Springer-Verlag Lecture Notes in Computer Science, 1996.

[16] J. R. Shewchuk. What is a good linear finite element? - interpolation, conditioning, anisotropy, and quality measures. Technical report, In Proc. of the 11th International Meshing Roundtable, 2002.

[17] S. M. Shontz and P. Knupp. The effect of vertex reordering on 2d local mesh optimization efficiency. In R. Garimella, editor, *Proceedings of the 17th International Meshing Roundtable*, pages 107–124. Springer Berlin Heidelberg, 2008.

[18] M. M. Strout and P. D. Hovland. Metrics and models for reordering transformations. In *Proceedings of the 2004 Workshop on Memory System Performance*, pages 23–34. ACM, 2004.

[19] P. H. T. Munson. The feasnewt benchmark. In *Proceedings of the IEEE International Workload Characterization Symposium*, pages 150–154, 2005.

[20] M. Zhou, O. Sahni, M. S. Shephard, C. D. Carothers, and K. E. Jansen. Adjacency-based data reordering algorithm for acceleration of finite element computations. *Sci. Program.*, 18(2):107–123, Apr. 2010.

		Quantiles				#accesses
mesh	Ordering	50%	75%	90%	100%	
carabiner	ORI	8	52	1,168	1,924,021	15,566,520
	BFS	1	11	99	1,923,989	
	RDR	1	4	6	1,942	
crake	ORI	8	43	642	1,767,468	14,226,264
	BFS	1	11	80	1,767,488	
	RDR	1	4	6	3,903	
dialog	ORI	7	39	306	1,819,234	14,614,336
	BFS	1	10	79	1,803,850	
	RDR	1	5	11	6,198	
lake	ORI	7	38	270	2,224,176	17,850,952
	BFS	1	10	72	2,224,069	
	RDR	1	4	5	8,774	
riverflow	ORI	7	37	1,430	1,963,783	15,758,200
	BFS	1	10	69	1,969,058	
	RDR	1	4	6	5,429	
ocean	ORI	7	38	3,866	2,312,947	18,719,512
	BFS	1	10	76	2,339,305	
	RDR	1	5	9	6,122	
stress	ORI	7	44	358	1,817,752	14,862,888
	BFS	1	10	91	1,817,758	
	RDR	1	4	6	4,809	
valve	ORI	7	40	646	1,780,959	14,301,320
	BFS	1	10	75	1,780,916	
	RDR	1	4	6	7,256	
wrench	ORI	7	42	376	2,290,783	18,429,528
	BFS	1	10	89	2,290,734	
	RDR	1	4	6	6,487	

Table 2: Distribution of the reuse distances of the first iterations of each run as a function of the meshes and the orderings.